

two storage devices or at least two partitioned groups of storage devices; and wherein said managing comprises balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said plurality of viewers from said at least two storage devices or said at least two partitioned groups of storage devices.

5

31. The method of claim 30, further comprising monitoring said workload distribution across said at least two storage devices or at least two partitioned groups of storage devices.

10

32. The method of claim 30, wherein said managing comprises ensuring that an amount of data consumed by each of said viewers during a cycle time of said storage system is less than or equal to the product of number of data blocks fetched from said at least two storage devices or partitioned groups of storage devices during a service time of said at least two storage devices or partitioned groups of storage devices, and a size of said fetched data blocks.

33. The method of claim 30, wherein said managing comprises ensuring that total available buffer space of said storage system is greater than or equal to the product of number of data blocks fetched from said at least two storage devices or partitioned groups of storage devices during a service time of said at least two storage devices or partitioned groups of storage devices, and a size of said fetched data blocks.

34. The method of claim 30, wherein said managing comprises ensuring that an amount of data consumed by each of said viewers during a cycle time of said storage system is equal to the product of number of data blocks fetched from said at least two storage devices or partitioned groups of storage devices during a service time of said at least two storage devices or partitioned groups of storage devices, and a size of said fetched data blocks; and by ensuring that total available buffer space of said storage system is greater than or equal to the product of number of data blocks fetched from said at least two storage devices or partitioned groups of

storage devices during a service time of said at least two storage devices or groups of storage devices, and a size of said fetched data blocks.

5 35. The method of claim 30, wherein said buffer memory space comprises a part of an integrated cache/buffer memory of said storage system.

10 36. The method of claim 35, further comprising monitoring a number of viewers that are reading data from said at least two storage devices or partitioned groups of storage devices out of the total number of viewers being served by said storage system; and wherein said managing comprises balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said viewers reading data from said at least two storage devices or partitioned groups of storage devices.

15 37. The method of claim 35, wherein said method further comprises limiting a size of said buffer memory space by implementing a read-ahead buffer limit.

20 38. The method of claim 30, wherein said managing comprises setting a cycle time of said at least two storage devices or partitioned groups of storage devices to be greater than or equal to the maximal aggregate service time of said at least two storage devices or partitioned groups of storage devices.

25 39. The method of claim 38, wherein said managing further comprises setting a cycle time of said at least two storage devices or partitioned groups of storage devices to maximize the number of simultaneous viewers of said continuous media data that is supported by said information management system.

30

40. The method of claim 30, wherein said method further comprises allocating said I/O resources between background processing activities and delivery of said continuous media data.

5

41. The method of claim 30, wherein said managing comprises at least one of performing I/O admission control, determining read-ahead size, or a combination thereof.

10 42. The method of claim 41, wherein said managing comprises performing said I/O admission control by monitoring the number of existing viewers served from said at least two storage devices or partitioned groups of storage devices, and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers; and determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space.

74

15
20 43. The method of claim 41, wherein said managing comprises determining said read-ahead size by monitoring the number of existing viewers served from said at least two storage devices or partitioned groups of storage devices, and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers; setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, said monitored data consumption rate, and a size of said data blocks.

30